RSAC 2025 Conference

Many Voices.
One Community.

SESSION ID: MASH-W02

AI Is Just Math: Get Over It!

Ira Winkler, CISSP

CISO
CYE Security
https://www.linkedin.com/in/irawinkler/

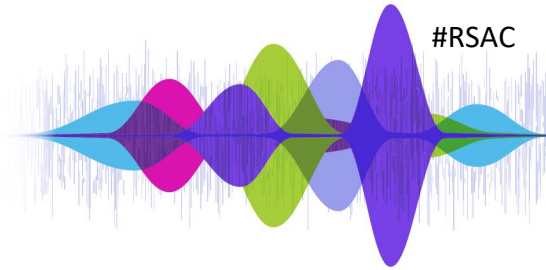

Disclaimer

Presentations are intended for educational purposes only and do not replace independent professional judgment. Statements of fact and opinions expressed are those of the presenters individually and, unless expressly stated to the contrary, are not the opinion or position of RSA Conference LLC or any other co-sponsors. RSA Conference LLC does not endorse or approve, and assumes no responsibility for, the content, accuracy or completeness of the information presented.

Attendees should note that sessions may be audio- or video-recorded and may be published in various media, including print, audio and video formats without further notice. The presentation template and any media capture are subject to copyright protection.

© 2025 RSA Conference LLC or its affiliates. The RSAC and RSAC CONFERENCE logos and other trademarks are proprietary. All rights reserved.

Where We Are Apparently Heading

Where We Are Apparently Now

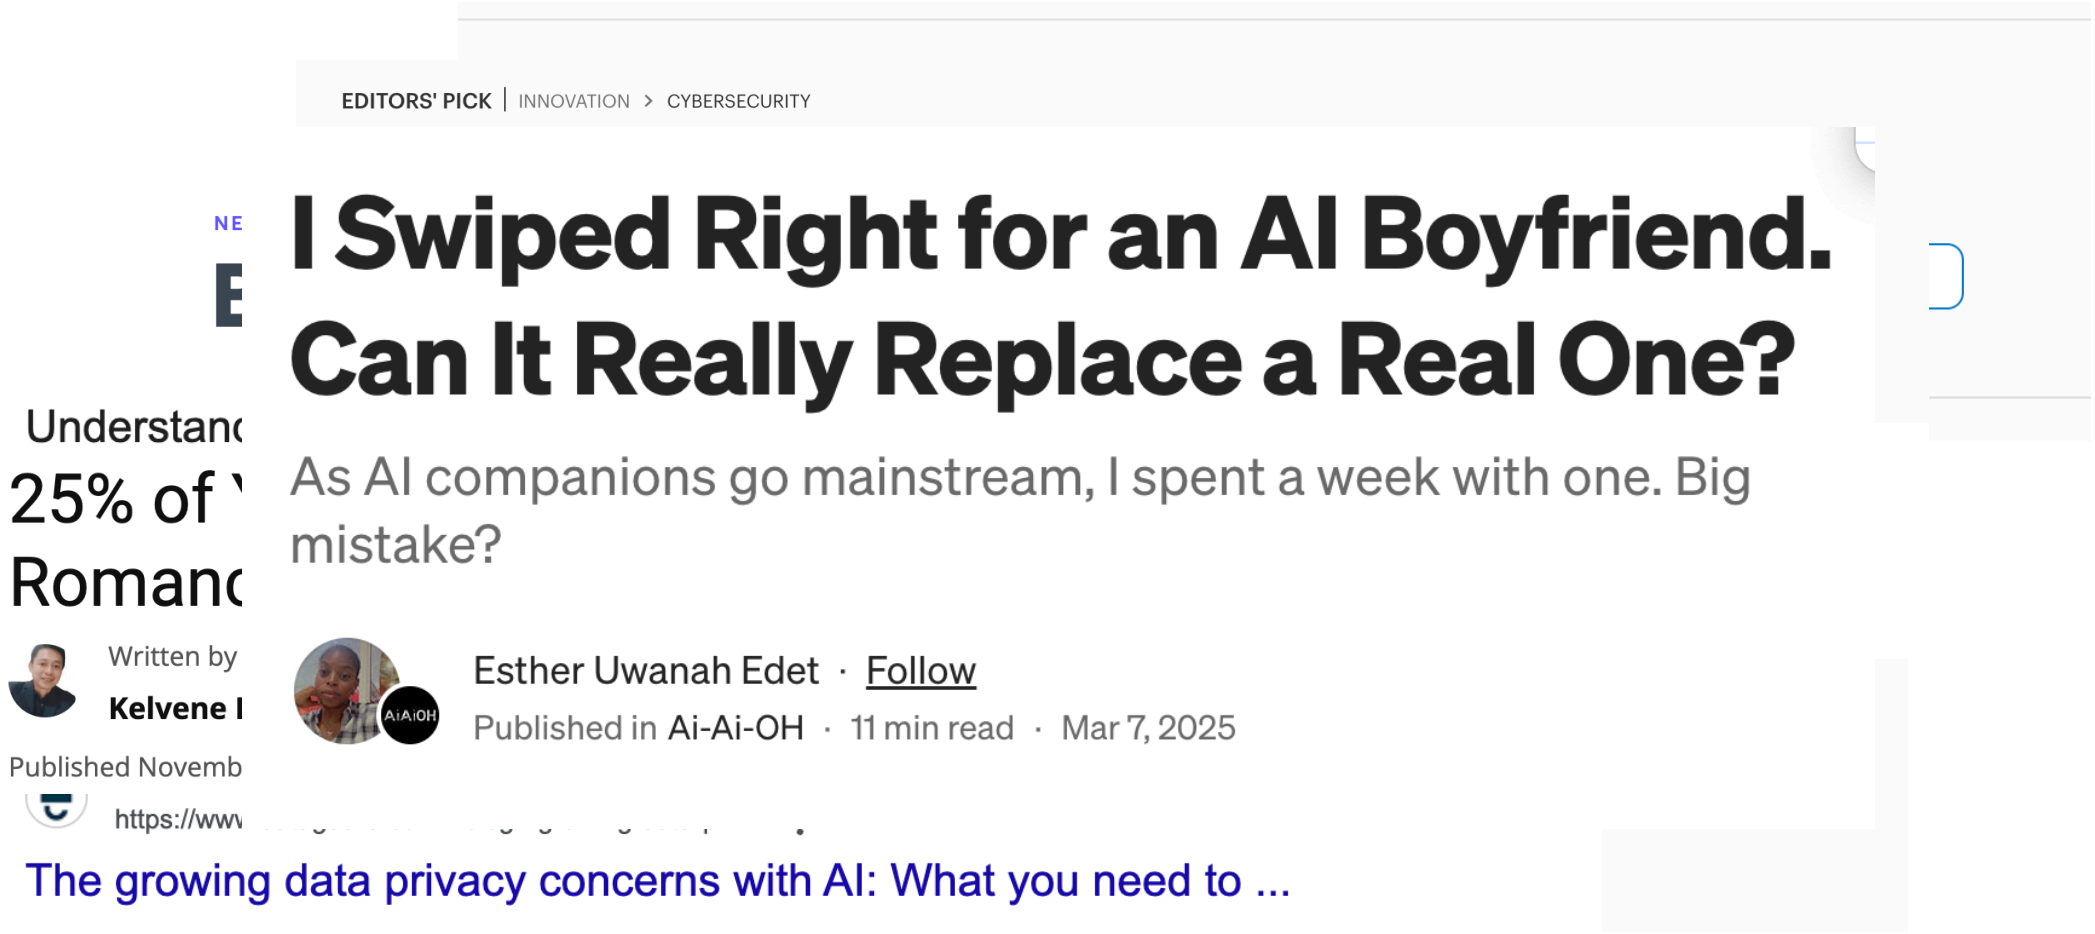

EDITORS' PICK | INNOVATION > CYBERSECURITY

I Swiped Right for an AI Boyfriend. Can It Really Replace a Real One?

As AI companions go mainstream, I spent a week with one. Big mistake?

Esther Uwanah Edet · Follow

Published in Ai-Ai-OH · 11 min read · Mar 7, 2025

Understanc

25% of

Romanc

Written by

Kelvene

Published Novemb

https://www

The growing data privacy concerns with AI: What you need to ...

Practical Applications

RSAC 2025 Conference

The Ira-ness of This Presentation

90%+ of “AI experts” have no clue what they’re talking about

Many Voices.
One Community.

RSAC 2025 Conference

Even the 10% Know Very Little

But that's OK.

Think medical specialties.

Many Voices.
One Community.

What Exactly Is AI?

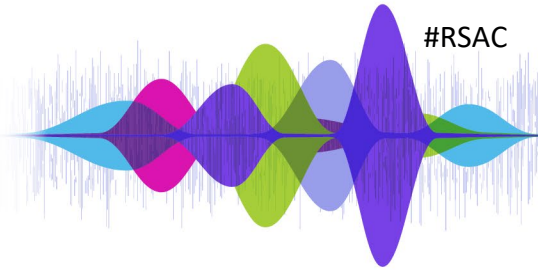

Artificial intelligence is a field of science concerned with building computers and machines that can reason, learn, and act in such a way that would normally require human intelligence or that involves data whose scale exceeds what humans can analyze.

A Few Points

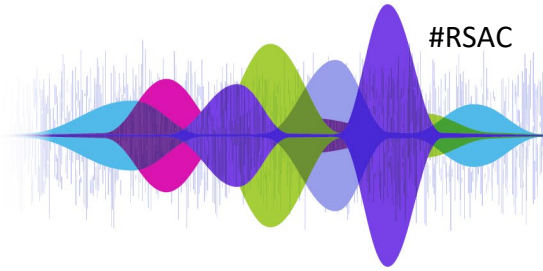

- The same can be said for the advent of computers as a whole
- Computers and machines are primarily software algorithms embedded in computers
  - Cars, robots, etc all have embedded computers running the software
- Chips might have embedded firmware
- It's again all software algorithms that have frequently been around for decades

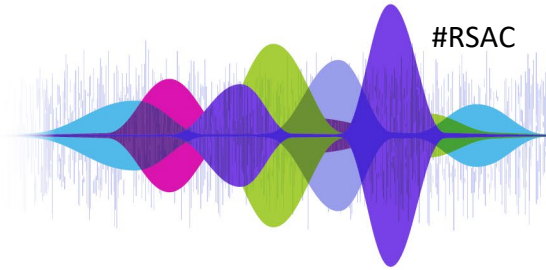

Most Importantly

- AI is not a singular entity
- Anyone generically using the term, “AI”, is doing you a disservice

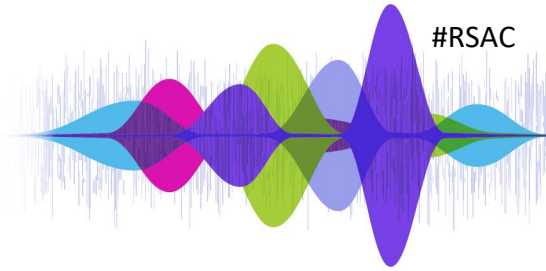

Why Now?

- Algorithms require a lot of data for training
- Require a lot of data to process for best results
- Require a lot of processing capability
  - Why NVIDIA is now worth trillions of dollars
  - Why DeepSeek tanked NVIDIA
    - Demonstrated algorithms can be more efficient with processing needs

We've Been Using AI For Decades

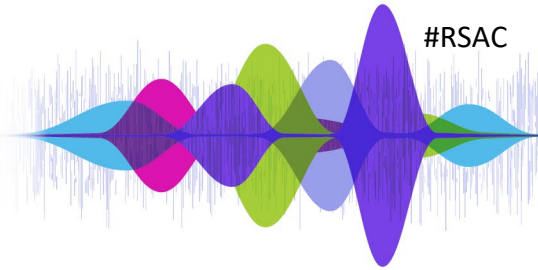

- Anti-malware
- OCR
- Siri, Alexa, etc
- Voice recognition
- CGI
- Countless mundane algorithms

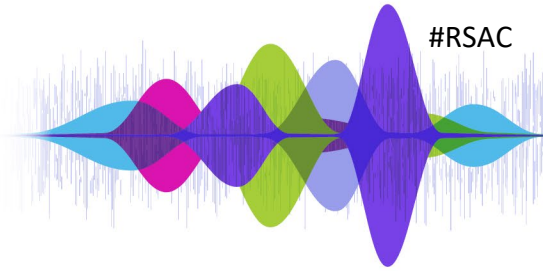

Mortgage Processing

- Manual
  - Basic credit scores and logic
- Statistics
  - Statistical analysis of zip codes
- AI (Machine Learning)
  - Decision trees, regression analysis, etc
  - More factors taken into account
  - Hopefully providing better results

How Did Target Know A Teen Was Pregnant?

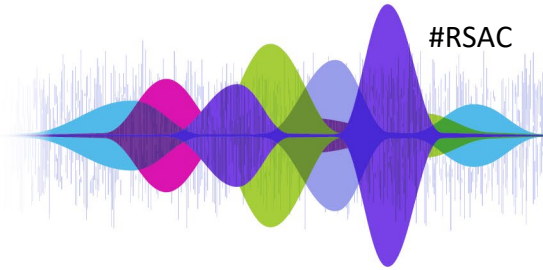

- It didn’t
- Father outraged by advertisements sent to his daughter
- Target used data mining to determine buying patterns
- Teen exhibited a “buying pattern”
- Data mining determined that people exhibiting the buying pattern would enter another buying pattern in a period of time
- Advertisements for sponsored SKUs sent in advance of predicted change in buying pattern
- They never knew or cared the teen was pregnant, but wanted to promote SKUs

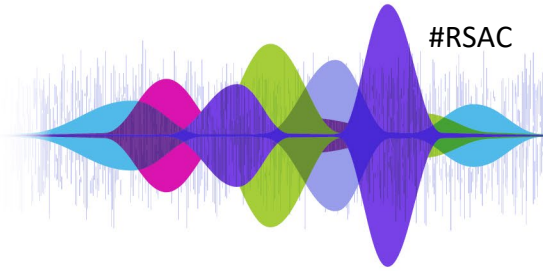

Other Examples

- Cambridge Analytica
- Facebook, Twitter, Instagram algorithms
- TikTok, Netflix, and others reportedly determining if women are bisexual
- Credit targeting
- Netflix and Amazon recommendation engines
- Predictive text
  - You've got to be ducking kidding me
- Loan approvals

Why Understanding What AI Actually Is Is Important

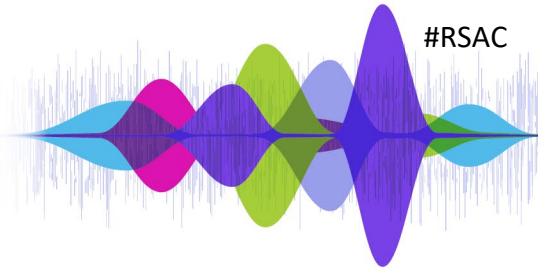

- When everything is AI, nothing is AI
- Clumping all AI together is problematic
- Can one law properly regulate computer vision, data mining, robots, self-driving cars, deep fakes, LLMs?
- Fundamentally, it is an implementation of diverse computer algorithms that is poorly defined

AI Requirements
(An oversimplification)

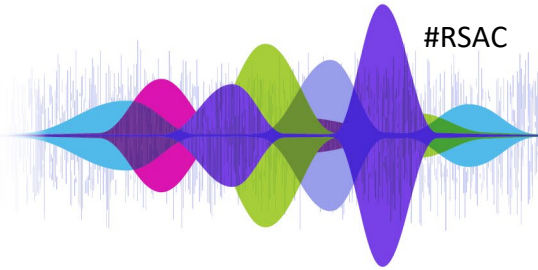

- An algorithm
- Identify relevant attributes
- Training data
- Trial and error with attribute choice and weighting

Results Can Change

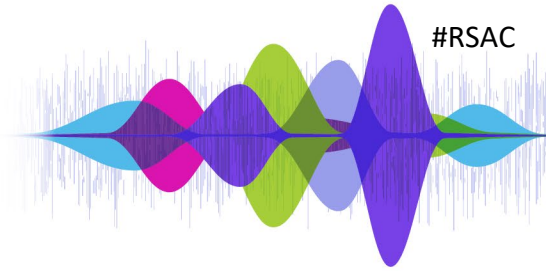

- A key aspect of the algorithms is that output will change as more ”training data” becomes available
- Additional training data means more refined results
- Output results can modify weightings
- Consider Netflix algorithms improving when you provide more data

AI Isn't Biased, The Programming Is Incomplete

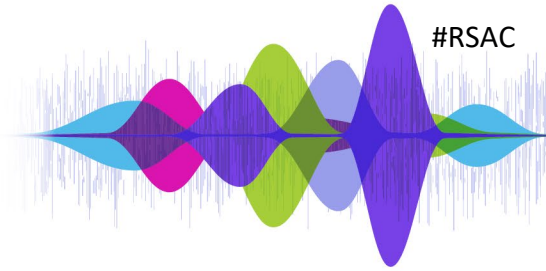

- Misdiagnosis of African Americans in medical diagnostic systems
  - Underdiagnosed hypertension
  - Root problem was the data scientists did not want to seem racist and didn't use race as an attribute in early models
  - Added race and gender as attributes for advanced model
- Self-driving cars crash
  - Need to understand why to improve the implementation
- The question is, "When is a model good enough for its intended use?"

Determining Attack Paths And Optimal Countermeasures

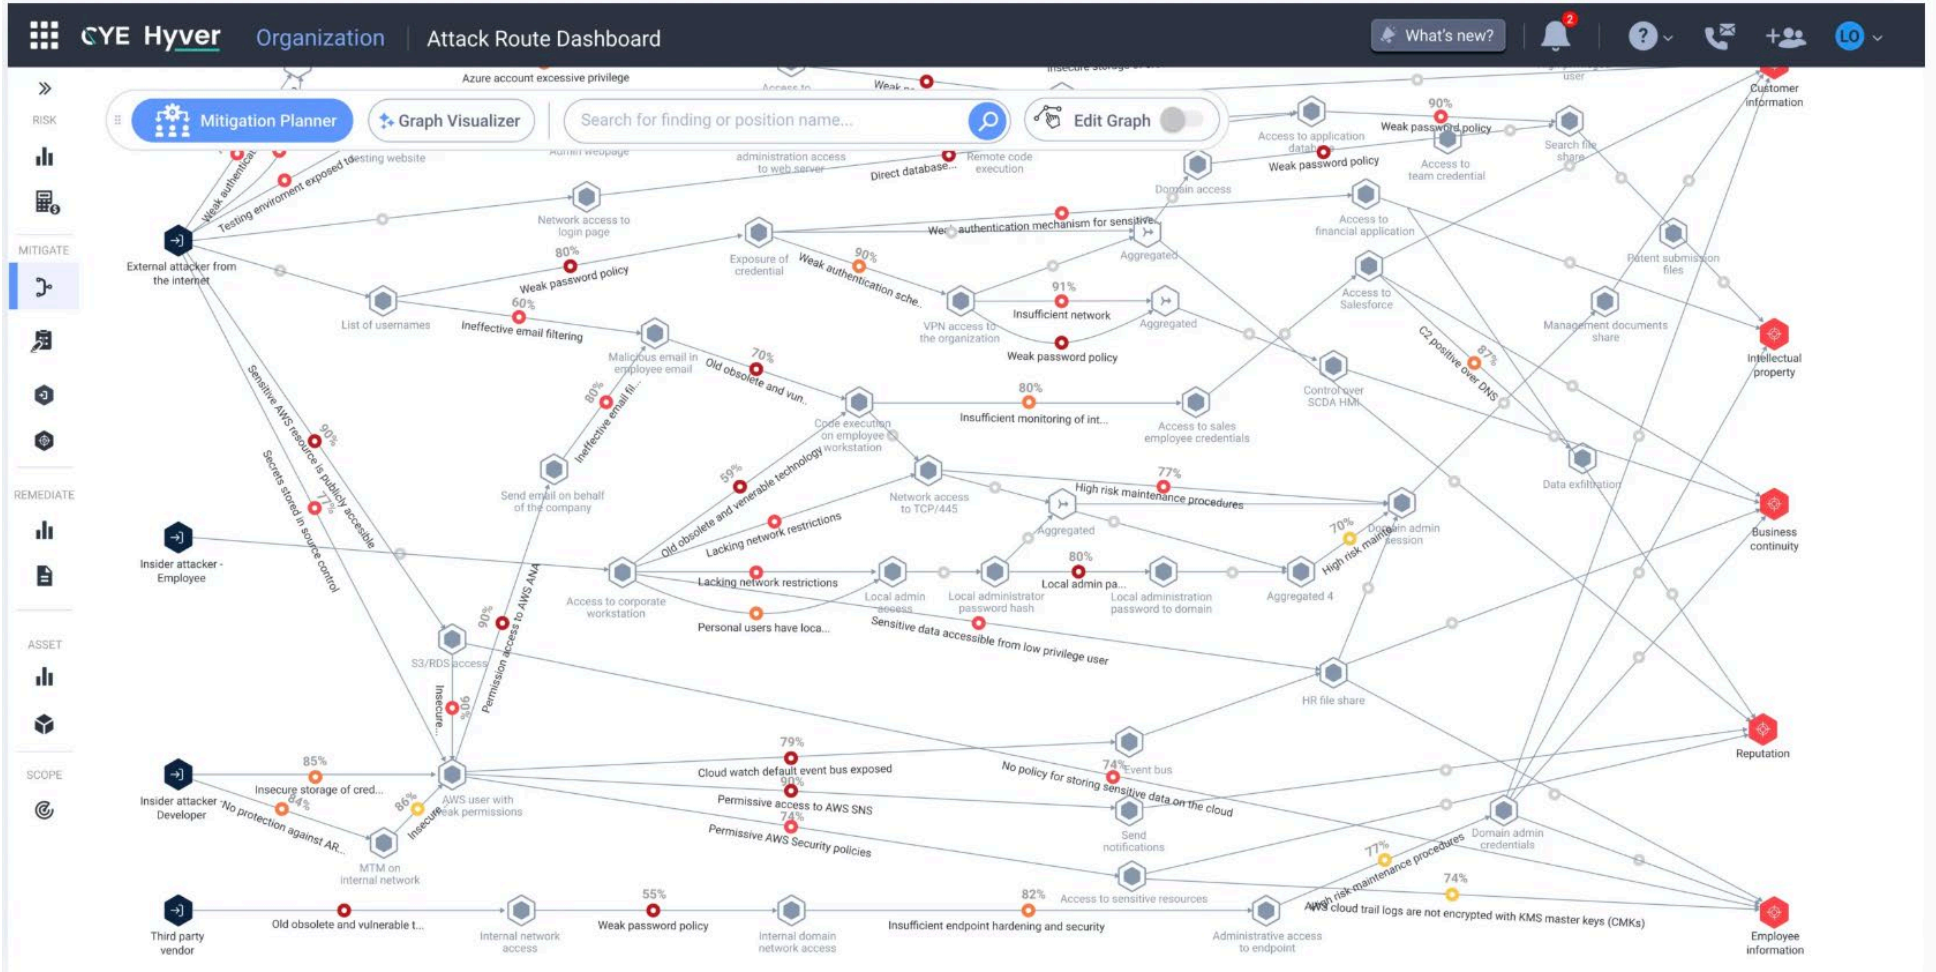

RSAC 2025 Conference

But This Is Real AI!!!

Many Voices.
One Community.

Automated Access Provisioning

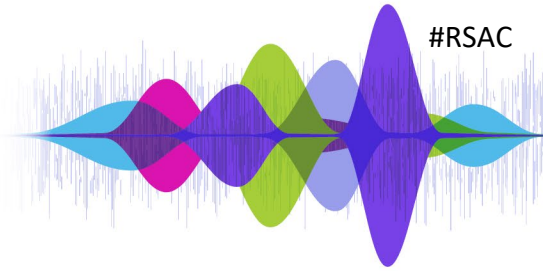

- User requests access to financial database
- User never accessed data before
- Similar requests denied or time-limited
- Teammates have access to 4 tables
- Request is for a short-term project
- Read-Only access for 7 days granted

Is It Really That Much More Advanced?

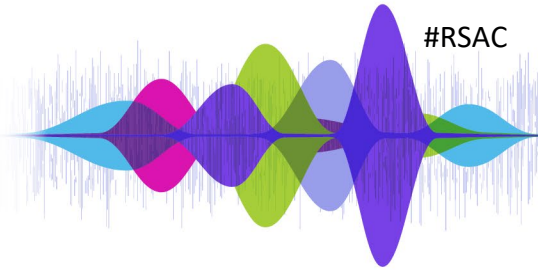

- It does provide context more than deterministic
- It appears to think things through
- It chains decisions together
- It determines new questions to ask
- It goes through data in ways humans could put is time consuming
  - That is what computers do
- Is it that much different than other applications?
- Pretty similar to LLM concepts, but less involved

What Is Accelerating “AI”?

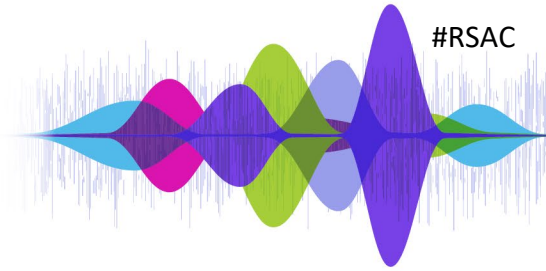

- It really is adoption of the algorithms
- Requires understanding of advanced mathematics
- Requires data availability
- Requires processing capability
- Requires considering, what if?

Supervised vs Unsupervised Learning

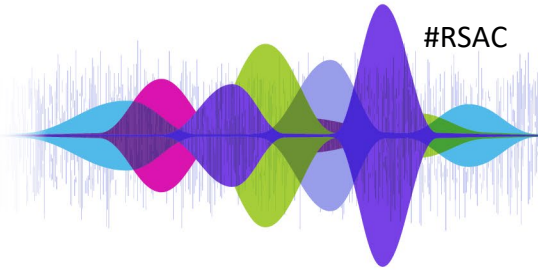

- A major categorization of Machine Learning
- Unsupervised Learning
  - Exploratory
  - Pattern searching
    - Target pregnant teen
    - Netflix and Amazon recommendations
    - TikTok feeding women content targeted to bisexuals
- Supervised Learning
  - Know the answer
  - Determining probability of finding the answer

Supervised Learning Cycle

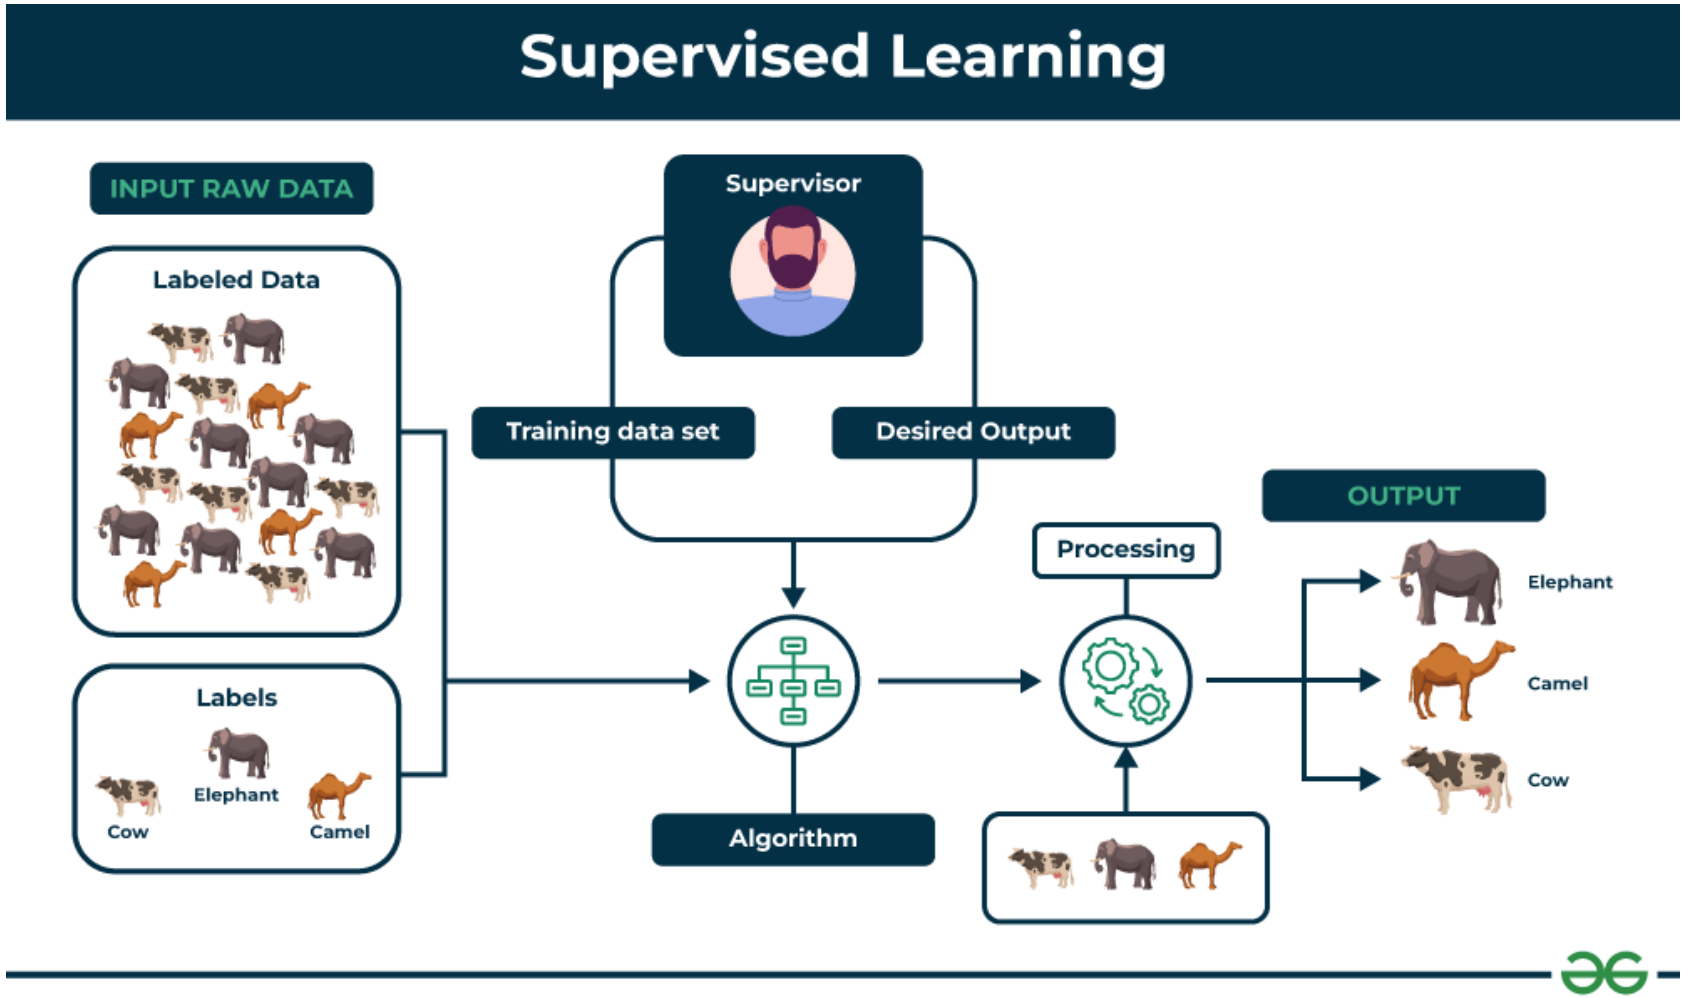

Example Of Inexact Decision

MNIST 0-9 Kaggle A-Z

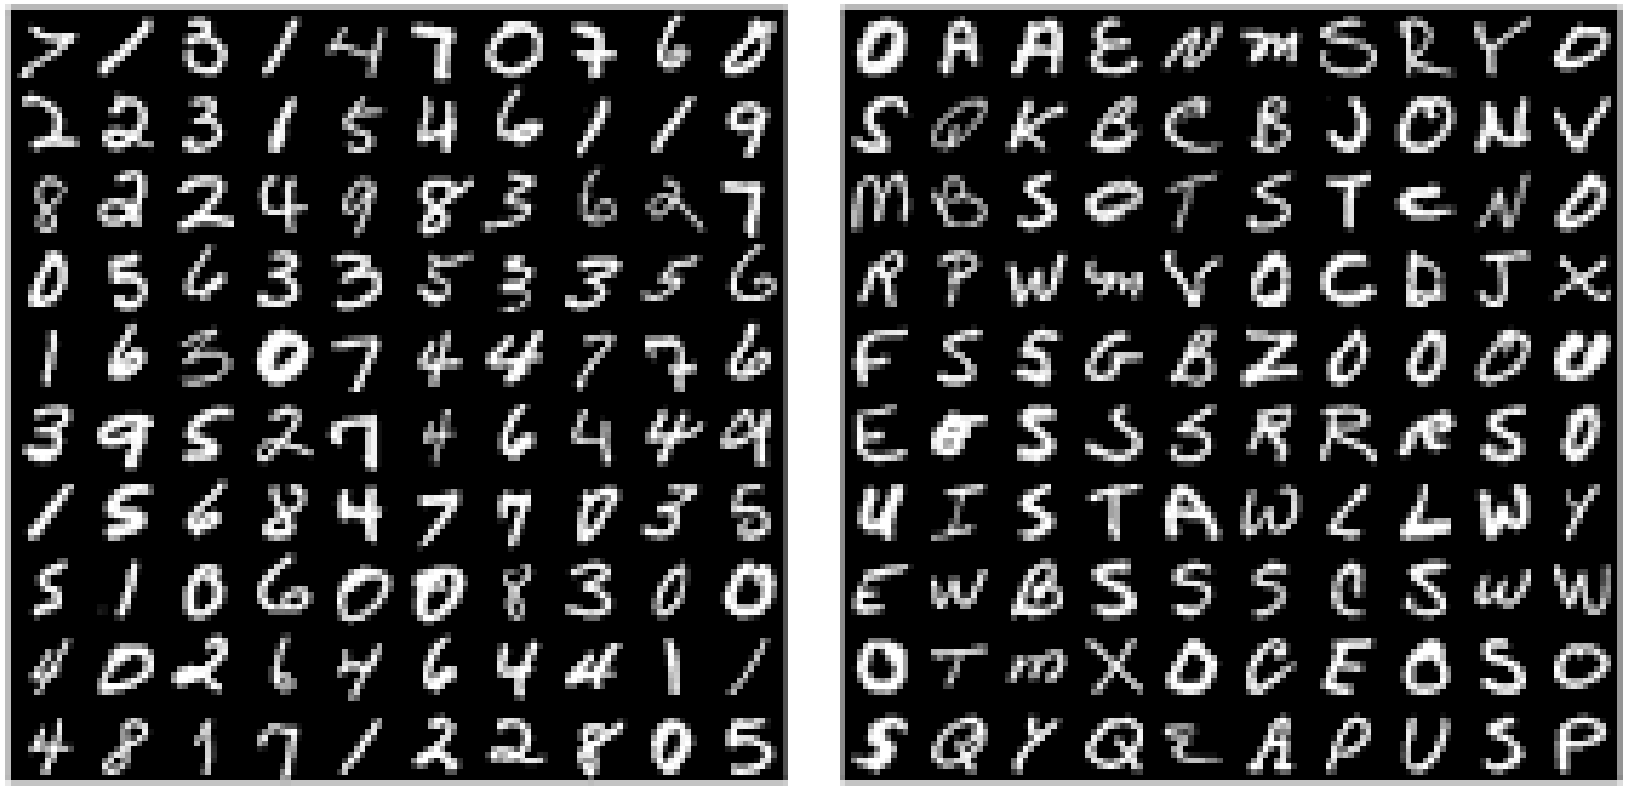

Sample Math

$$P\left(B_J \mid A\right) = \frac{P\left(A \mid B_J\right) P\left(B_J\right)}{\sum_{i=1}^{n} P\left(A \mid B_i\right) P\left(B_i\right)}$$

$$\operatorname{simil}(x, y) = \frac{\sum_{i \in I_{xy}} \left(r_{x,i} - \bar{r_x}\right)\left(r_{y,i} - \bar{r_y}\right)}{\sqrt{\sum_{i \in I_{xy}} \left(r_{x,i} - \bar{r_x}\right)^2} \sqrt{\sum_{i \in I_{xy}} \left(r_{y,i} - \bar{r_y}\right)^2}}$$

$$\arg\min_{\mathbf{S}} \sum_{i=1}^{k} \sum_{\mathbf{x} \in S_i} \left\| \mathbf{x} - \boldsymbol{\mu}_i \right\|^2 = \arg\min_{\mathbf{S}} \sum_{i=1}^{k} |S_i| \operatorname{Var} S_i$$

Example For Phishing Susceptibility

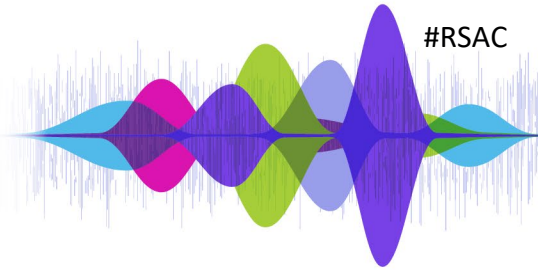

- Had more than 900 attributes on 130 users
- Performed ANOVA (traditional statistics) on key attributes
  - Close, but no significance
- Attempted ML/Data Mining techniques
  - Randomly tried different algorithms with different attributes
  - Nothing better than random groupings
- After more than 100 tests, K-Means Cluster Analysis, with 6 clusters, on 5 attributes
  - 100% accuracy in identifying phishing susceptibility

Even The Most Advanced Algorithms Have These Issues At Their Core

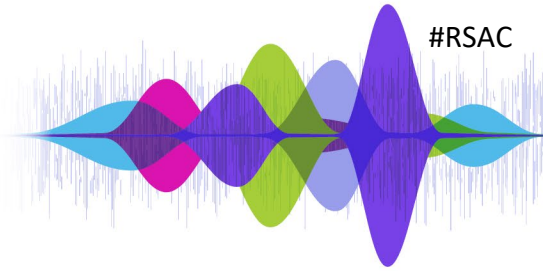

- Robots need to categorize what they're looking at
- Autonomous vehicles have the same concerns of categorizing their environmental
- Military applications need to understand friend or foe
- Deep fake detection needs to categorize voice and video in real-time

RSAC 2025 Conference

The Advance That Advances Things

Automating the choice of algorithms, attributes, and tuning

Many Voices.
One Community.

Requirements For “AI” Applications

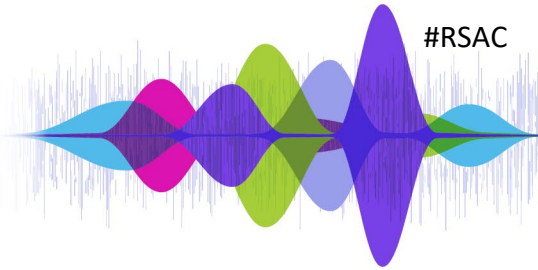

- Data sets/Training data
  - Is it a good set of data? Can someone poison it? Is it constantly updating?
- Model
  - Which algorithm? How can you tune it? Can you hack it?
  - Is it just lucky? Retuning the model
- Data output
  - What are you doing with data? How critical is the output?
  - Is the output protected? Are you protecting the output?

These Are Fundamental Concepts

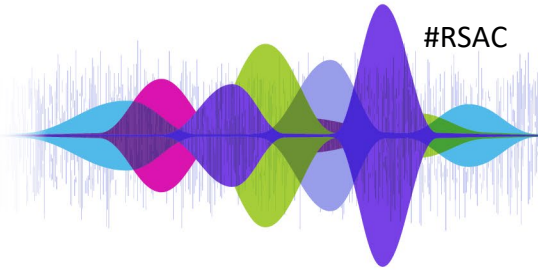

- Applies to self-driving cars, OCR, computer vision, Netflix, mortgages, anti-malware, credential allocation, deep fakes, etc
- Yes, there are differences between many applications
- Focus on the basics to reduce overwhelm
- Understand that implementing AI means determining a problem and testing a variety of algorithms against a problem
- It could be wrong. It could just need tweaking
  - 6 clusters vs 5 or 7 clusters

AI Concerns

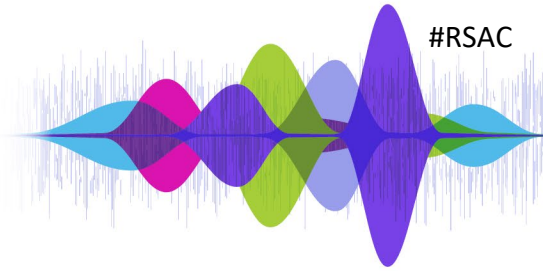

- Privacy
  - Input data? Implications of output?
- Security
  - Input? Output? Autonomy? Human intervention? Trust? Testing? Training data?
- Deep fakes, public manipulation, etc
  - Complicated, but requires a combination of processes, regulation, and technologies

Will “AI” Take Your Job?

- Same was said about computers
- Perhaps it will enhance your capability
- If your job can be boiled down to an algorithm, you’re at risk
- Manual labor can be automated
- The difficulty of the algorithms required will drive the adoption
- If it can be done, someone will do it
  - Theoretically, we need to develop boundaries proactively

Every Application Has Its Unique Concerns

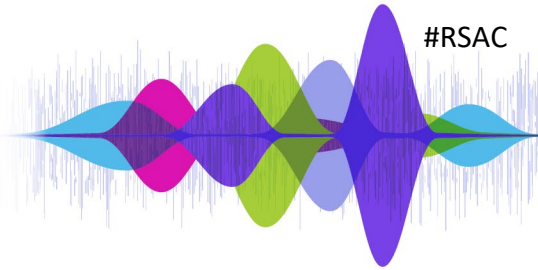

- Stop referring to it as ”AI”
- It’s as ubiquitous as computers
- Address the concerns of the specific applications
- This means the ethics, legality, and implications of the input and output to an application, regardless of the algorithms in use
- If you’re talking agentic AI acting autonomously, say it specifically
  - And with the required context
- Please stop the underlying hype and ignorance

Input, Process, Output

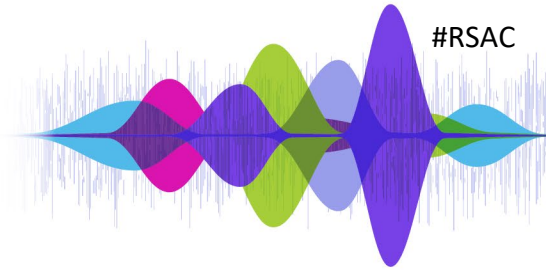

- Compliance, governance, regulation, etc. must be specific to the application
- There are commonalities through all, but they are basically math
- You can’t regulate math, but you can regulate the input and output
  - Every application needs its own regulation to work
- Force the “experts” to talk context, not vagaries

Results Can Be Invasive

- Math or not, output must be evaluated
- Pattern detection is incredibly powerful
- Knows things about people they don't know
- Not specific to "AI", but more prominent given the nature of the algorithms
- The same, but worse

RSAC™ 2025 Conference

Chaining Decisions Is The Direction Things Are Going In For Advances

Apply What You Have Learned

- Determine the specific field(s) of AI that are relevant to your organization, and to you as a professional
- Seek out understanding of the field(s)
- Determine the concerns of the inputs and outputs of your relevant applications
- Determine if you to examine the veracity of the algorithms and how much involvement you need to have
- Create and implement an action plan
- Determine the need to secure the process, inputs, and outputs and act accordingly
- Be skeptical of anyone talking about AI in generalities versus their specific expertise

RSAC™ 2025 Conference

Many Voices.
One Community.

Ira Winkler

ira@cyesec.com

www.cyesec.com